

# The InSight APSS Data Return Anomaly: Development of an Automated Detection and Response Method

Emily Manor-Chapman, Elizabeth Barrett, Farah Alibay, Kyle Cloutier,  
Jonathan Grinblat, Jesse Mendoza Jr, Nimisha Mittal  
Jet Propulsion Laboratory, California Institute of Technology  
4800 Oak Grove Drive  
Pasadena, CA 91109  
emily.a.manor@jpl.nasa.gov

**Abstract**— The Auxiliary Payload Sensor Suite (APSS), a collection of environmental sensors carried by the Interior exploration using Seismic Investigations, Geodesy, and Heat Transport (InSight) lander, is capable of measuring Martian air temperature, wind speed, atmospheric pressure, and local magnetic fields. After beginning Mars surface operations, the instrument experienced an anomaly that prevented it from returning science data. The anomaly affected not only the instrument, but also had impacts at the system level. APSS returned to normal operations, however the anomaly occurred again just several weeks later. This proved the need for a streamlined recovery response that would be adaptable to the operations planning cycle and workforce, that would limit the system-level impacts of the anomaly, and that would minimize the instrument downtime. The recovery response evolved from a ground-in-the-loop response to an onboard method for detecting occurrences of the anomaly and automatically recovering the instrument. Ultimately, the automated detection and response method reduced instrument downtime from days to hours and significantly minimized science data loss.

## TABLE OF CONTENTS

1. INTRODUCTION .....	1
2. APSS INSTRUMENT .....	2
3. DATA RETURN ANOMALY .....	4
4. RECOVERY PROCESS .....	6
5. LESSONS LEARNED .....	12
6. SUMMARY .....	12
ACKNOWLEDGEMENTS.....	12
REFERENCES .....	12
BIOGRAPHY .....	13

## 1. INTRODUCTION

The InSight mission is the first mission to focus on Mars' interior structure and evolution by conducting seismic, heat transport and geodesy investigations [1]. The data collected by InSight will provide information on the formation and processes of rocky planets. The mission launched from Vandenberg Air Force Base on May 5, 2018. After a six-month cruise, the lander arrived at Mars on November 26, 2018, landing in Elysium Planitia. The prime science mission will last for one Martian year (approximately two Earth years).

## Lander Overview

The lander, provided by Lockheed Martin Space, is based on the design of NASA's Mars Phoenix Lander. It is solar powered and is responsible for telecommunications, command and data handling (C&DH), power, and thermal control. The lander structure includes the deck, which is the top side of the lander where components such as the communication antennas and payloads are located, and the thermal enclosure which is located under the deck and houses the lander avionics and payload electronics. The solar arrays extend like circular wings on either side of the lander deck.

The payloads carried by the lander include:

- Seismic Experiment for Interior Structure (SEIS), built and operated by Centre National d'Études Spatiales (CNES) and their partners, is a seismometer that continuously monitors seismic activity via measurements of the surface ground velocity [2].
- APSS, described further in Section 2, provides environmental context for the mission's seismic investigation and aims to return one of the most complete weather datasets from the red planet via high-rate, continuous sampling [3].
- The Heat Flow and Physical Properties Probe (HP<sup>3</sup>), built and operated by Deutsches Zentrum für Luft- und Raumfahrt (DLR), is designed to determine the geothermal heat flux and surface brightness temperatures of Mars [4].
- The geodesy investigation, Rotation and Interior Structure Experiment (RISE), utilizes the lander X-band telecommunications subsystem to measure planetary rotational variations [5].
- The Instrument Deployment System (IDS), built and operated by the Jet Propulsion Laboratory (JPL), consists of a robotic arm and two cameras (one attached to the arm and the other underneath the lander deck) [6]. The arm was used to place SEIS and HP<sup>3</sup> onto the Martian surface.

The lander and payloads are shown in Figure 1.

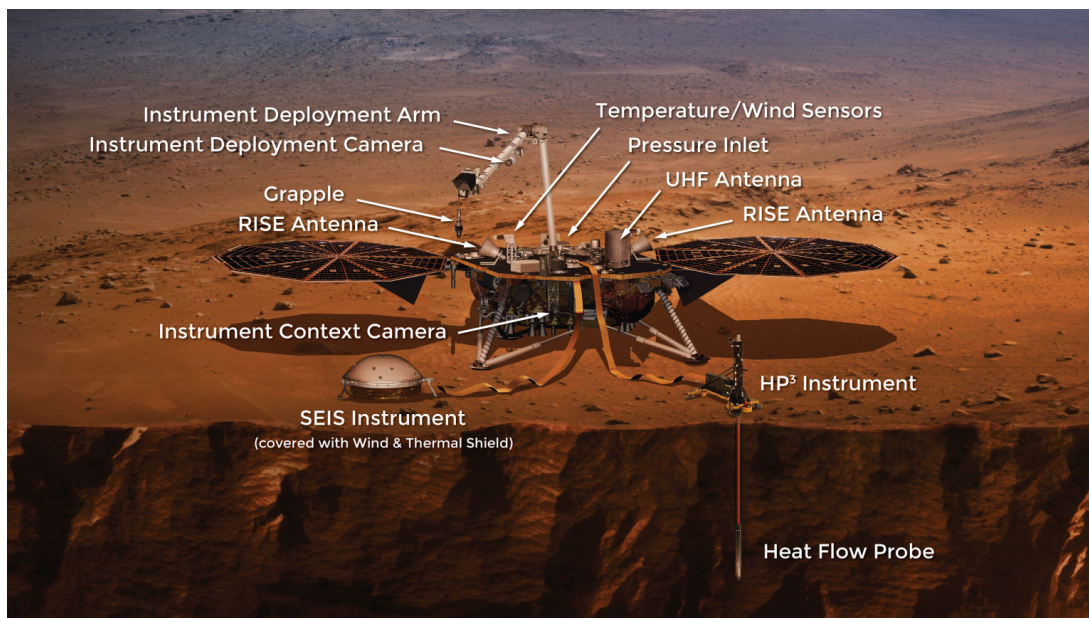
As the lander is solar powered, the lander cannot be continuously operated throughout a sol (Martian day); a sleep/wake cycle is used to manage power resources. During a wake period, all lander subsystems, such as communications and C&DH, can operate, but while the lander is asleep these functions are unavailable. However, the instruments are designed to operate throughout a lander sleep cycle with power, but no command or telemetry interface, from the lander. This allows science data to be collected continuously.

InSight uses two types of lander wake cycles: full wakeups and diagnostic wakeups. Full lander wakeups are longer in duration and allow for science and engineering data to be transferred from the instruments to the lander and for any planned instrument commands to be transmitted from the lander C&DH to the respective instrument. Shorter diagnostic lander wakeups are used to limit the duration of sleep cycles and periodically check in on the state of the system. Only engineering data is transferred between the instruments and lander during these wakes. Insufficient time is available to perform any other instrument activities during a diagnostic lander wakeup.

## Surface Operations

Surface operations are divided into two phases: deployment and science monitoring [7]. The deployment phase focused on selecting appropriate surface sites for the SEIS and HP<sup>3</sup> instruments and then placing them in their designated locations. Deployment was a high-activity period with near-daily tactical operations shifts. This involved staffing a large team of scientists and engineers to analyze downlinked data, plan for the next sol, and prepare the associated uplink products. This phase lasted approximately 100 sols.

The science monitoring phase began in Spring 2019, after both instruments were successfully placed on the surface and the commissioning of the SEIS instrument was complete. This phase will last for one Martian year. Operations staffing is substantially reduced in this monitoring phase with activity planning reduced from a daily to weekly cadence. With less staffing and a longer planning horizon, a lengthier timeframe is needed to respond to unexpected events.



**Figure 1: InSight Lander (Credit: NASA/JPL-Caltech)**

## 2. APSS INSTRUMENT


APSS, shown in Figure 2, is a collection of environmental sensors connected to a central electronics box, the Payload Auxiliary Electronics (PAE). The pressure subsystem (PS) consists of a pressure inlet (located on the top deck of the lander) and connector tube and a pressure sensor (located inside the lander thermal enclosure). This sensor measures atmospheric pressure. The pressure inlet and connector tube were provided by JPL. The pressure sensor was provided by Tavis Corporation.

The Temperature and Winds for InSight (TWINS) sensors, provided by Centro de Astrobiologia (CAB), measure air temperature and wind speeds. Two independent sensors are placed facing approximately opposite directions on the lander deck allowing winds from any direction to be accurately measured.

The InSight Flux Gate (IFG) magnetometer, provided by University of California Los Angeles (UCLA), measures local magnetic fields. IFG is located on the underside of the lander deck, outside the thermal enclosure.

The PAE (provided by JPL) controls power to and collects data from the sensors and interfaces with the lander and SEIS electronics box. It consists of a power board, digital board, pressure sensor board, magnetometer board, and a 28V monitor that measures lander bus voltage. The digital board contains the Field Programmable Gate Array (FPGA) and flash memory [8]. The FPGA governs the interfaces between the PAE and each sensor and the PAE and lander C&DH. It also controls memory management and science and

engineering packet generation. The flash memory is capable of holding up to 32 hours of science data packets. The flash memory read/write pointers are reset to the start of flash at each PAE power cycle, meaning that the locations of the pointers are not saved and cannot be restored [9]. Only science data packets are stored to flash memory. Engineering data packets are passed upon request to the lander C&DH in real-time during a lander wake cycle.



**Figure 2: APSS Instrument**

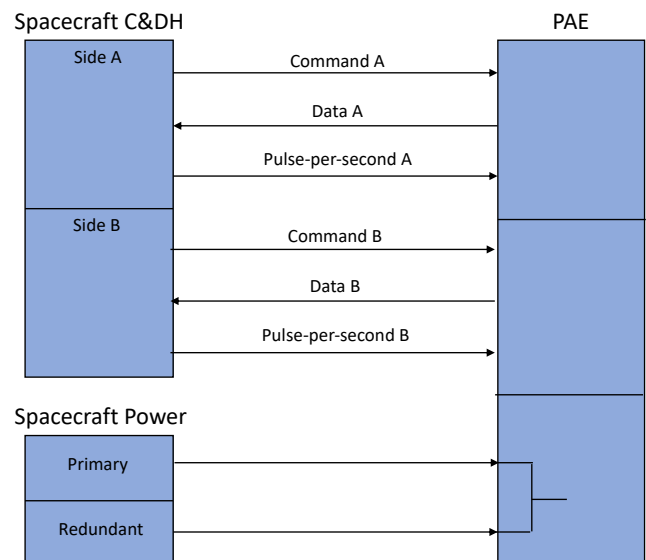
The PAE communicates with the lander, as shown in Figure 3, using a relatively simple electrical interface that is comprised of the following [10]:

- Redundant power from the lander
- Redundant low-speed, asynchronous RS422 for commanding from the lander to the PAE
- Redundant high-speed, synchronous RS422 for receiving data from the PAE to the lander
- Redundant pulse-per-second (PPS) via RS-422 used to keep time with an accuracy of at least four milliseconds

All communication is initiated by the lander and the PAE responds to commands. The PAE never sends data unsolicited to the lander. Only one “side” of each redundant signal is ever used at a time, and is dependent on which lander C&DH is active.

The lander flight software (FSW) includes a SEIS/APSS FSW module<sup>1</sup> that is responsible for processing all SEIS and APSS commands and data [11]. The PAE returns real-time engineering and science data to the lander. The lander can request science data in one of two ways: 1) as processed data or 2) as “raw” data packets. In nominal operations, science data is always requested as processed data. Processed data refers to data that has had finite impulse response (FIR) filtering and downsampling techniques applied in order to reduce the overall data volume that must be downlinked to

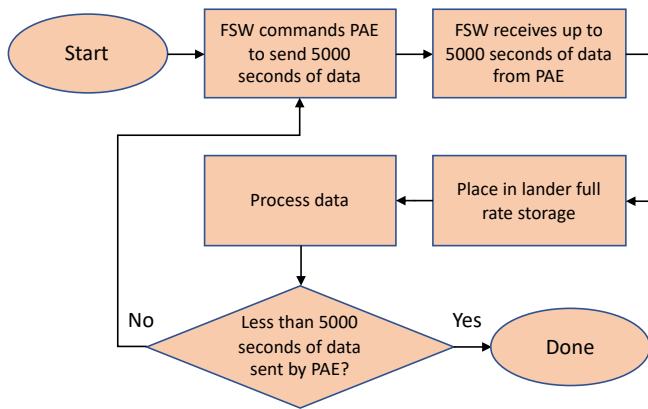
the ground. For example, TWINS data is collected at 1 Hz by the PAE, but can be downsampled to return a reading every 0.1 Hz. The filtering is done by the SEIS/APSS FSW. All data transmitted by the PAE to the lander is saved onboard in the full-rate storage, while only the downsampled data is queued for downlink. “Raw” data refers to packets that bypass the filtering process.



**Figure 3: PAE-Lander Interface**

<sup>1</sup> SEIS and APSS share an FSW module as both were developed by JPL and had similar software functional requirements.

When on-board sequences request science data processing from APSS, the steps occur as shown in the flow diagram in Figure 4.

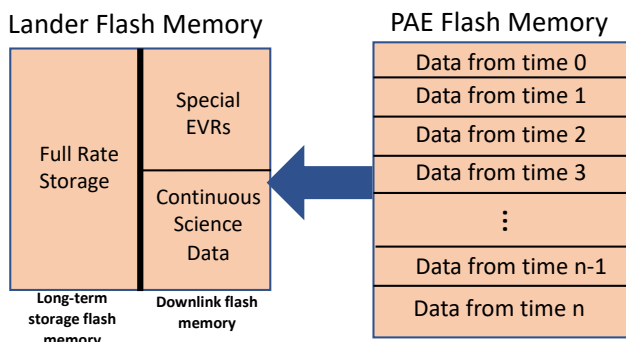


**Figure 4: APSS Data Processing Flow**

Flight software requests 5000 seconds of data at a time from the on-board storage of the PAE. With each 5000 second request, the raw science data is stored on the lander and simultaneously processed. If less than 5000 seconds of data is returned from the PAE in a given cycle, it means that the data transfer has reached the end of what is available in the PAE flash memory. As an example, if the PAE has been collecting data for five hours, then this will involve flight software doing four iterations of this process.

The “process data” step processes and funnels data to a number of different data types. Each one second set of data on the PAE is referred to as a packet. Each packet contains sensor power states for that second, error flags, timestamps, and samples of science data. For each packet of recorded data that the PAE sends, the lander FSW will split it up (see Figure 5) into the following types:

- Full rate science data
- Continuous science data
- Special Events Reports (EVRs)



**Figure 5: Organization of flash memory on the lander vs the PAE**

In the lander, the various types of data are stored in different flash memory locations. The continuous science data, as previously mentioned, contains a downsampled and/or filtered set of samples that covers the entire timespan of the recorded data from the PAE. All of this data is stored in the part of flash memory used to queue data for downlink to the ground.

The housekeeping data portion of each packet is processed into snapshots of state changes. Since the state of the PAE rarely changes once it is fully powered on and configured, this method saves on storage space and downlink bandwidth needed. These “snapshots” are referred to as Special Event Reports (EVRs). As an example, when a sensor power state changes from “on” to “off” as seen in the recorded data from the PAE, FSW will generate a Special EVR with the full state of the PAE at that timestamp. These Special EVRs are then sent to a separate buffer of the downlink queue in flash memory.

Lastly, the full set of science data is stored in the full rate storage. This data is compressed and organized by time. The full rate storage is designed to hold approximately 105 days’ worth of data, depending on the instrument configuration. This allows the science team time to analyze the continuous science data and then request the full, high resolution data set from a time period of interest.

### 3. DATA RETURN ANOMALY

The first surface operations activity for APSS was an instrument checkout on Sol 4. The checkout completed successfully, clearing the way for APSS to begin nominal operations on the following sol. On Sol 5, APSS powered on and began collecting data. However, after just a few hours, an anomaly occurred which affected science data collection and transfer. The anomaly caused the PAE to stop saving sensor readings to its flash memory. At the subsequent packet transfer from the PAE to the lander, the packets contained no science data and incorrect time tags. The instrument did continue to correctly report health and status via its real-time housekeeping (RTHK) engineering data packets. This data showed the instrument to be operating otherwise nominally. The operations team elected to safe<sup>2</sup> APSS while the anomaly was investigated further.

#### *Anomaly Symptoms*

The first indication in ground telemetry that an anomaly had occurred was an overflow of the Special EVR buffer. The buffer overflowed because the FSW interpreted the incorrect time tags in the science data packets as discontinuous times and thus created many time discontinuity reports.

The second indication was the receipt of a larger volume of continuous science data than expected. This overproduction

<sup>2</sup> To safe an instrument on InSight, the instrument is powered off and onboard flags are set to indicate a safed state. An instrument cannot return

to operations until those flags are cleared.

was due to inefficient compression of data during the science data processing session.

The third indication was the setting of the Flash Inhibit flag which is reported in the APSS RTHK packets. This flag indicates if writing of data to the PAE flash memory is enabled. Nominally, this flag should only be set if none of the instrument science sensors are powered on (as only science data is written to flash memory). However, if a science sensor is on and this flag is set, it indicates that a packet write has not occurred for at least two seconds (a write should occur once per second).

### *Initial Data Analysis*

The anomaly occurred during a lander sleep cycle, thus no lander telemetry, such as voltage and temperature data, nor APSS RTHK data was available to provide insight into the system at the time of interest. Also, without science data packets, the state of APSS at the time of the anomaly was unknown. Thus, initial data analysis had to rely on telemetry gathered at the surrounding lander wake cycles.

A review of the RTHK data gathered prior to and following the anomaly confirmed the following:

- All sensors are powered on, as expected.
- The time tags on the RTHK packets are as expected, confirming that the PAE did initialize time properly during power on.
- The 28V monitor was returning data (its readings are returned in both RTHK and science packets). This indicated that at least one sensor was producing readings correctly.
- Other PAE logic, such as that used to control the TWINS sensors, was functioning properly. This narrowed down the possible areas of the PAE that could be impacted by the anomaly.
- During the science data processing session, the PAE did not transfer as many packets as expected. A science data packet should be saved once per second to the PAE flash memory, however telemetry showed that the PAE transferred 3,261 less packets than it should have. From this data it was determined that writing to flash stopped at approximately 4.5 hours after instrument power on. This, combined with the lack of science data, confirmed that the anomaly involved the PAE science packet read and write functions.

A nominal science data processing session includes requesting a single raw packet from the PAE. This is done as a check of PAE responsiveness, in the event that the science data processing (Figure 4) fails to return any packets [12]. In this case, the raw packet also provided information as to the contents of the bad science packets. Inspection of the packet showed that only the header and checksum of the packet were correctly reported, while the contents of the packet was zero-

filled. Only the contents of the packet, that is the science data, are saved to the PAE flash memory. The packet header and checksum are created on-the-fly during transfer to the lander. Again, this pointed to an issue with the PAE flash read function.

Finally, lander telemetry did not show any errors while processing the science data packets. This illustrated that the format and the checksum of the packets were correct and the issue lay with the contents of the packets provided. If the packet format or checksum had been incorrect, the SEIS/APSS FSW would not have been able to process data.

### *Investigation*

After the initial data analysis, the investigation focused on three areas: ground testing, assessing instrument performance and retrieving additional data, and evaluating the anomaly cause.

*Ground Testing* – The Sol 5 activity plan was executed on the high-fidelity system testbed in an attempt to recreate the anomaly and gain additional data. Unfortunately, the anomalous behavior could not be replicated. However, this testing did allow for a comparison of the APSS RTHK data from flight and that generated by the testbed. The comparison confirmed that the PAE was responding to all commands and that it did underproduce science packets.

*Additional Data Retrieval* – While telemetry had shown when the anomaly occurred and pointed to a failure of the flash memory read and write functions, the operations team elected to retrieve data from the PAE flash memory in order to confirm what data had or had not been saved in memory prior to the anomaly. As discussed previously, the PAE does not track flash memory read and write pointers across power cycles. Therefore, to retrieve data from flash memory, the pointers must not be reset at power on and, instead, the read pointer must be kept at the start of memory and the write pointer at the end of memory. All data that is present in flash memory can then be read back via a standard science data processing session. This method was used on Sol 8 to retrieve all data that had been written to flash memory on Sol 5 (the anomaly sol).

A single raw packet was once again retrieved at the start of the science data processing session. This was the same raw packet as was retrieved on Sol 5. However, on the second retrieval the packet was populated as expected. Instead of being zero-filled, sensor readings were present. This indicated that science packets were being written correctly to flash memory until the anomaly occurred.

Also, the science data processing session completed nominally without any overproduction of data. The continuous science data contained sensor readings as expected, again confirming that the PAE flash write function was nominal until the time of the anomaly. Thus, whatever mechanism was causing the PAE to stop writing packets to



flash memory was also prohibiting the PAE from correctly reading previously created packets.

*Assessing Instrument Performance* – While the behavior was understood to be the result of a failure of the flash read/write functions, the question remained as to whether it was a transient or permanent failure. The operations team decided to repeat the APSS instrument checkout that had successfully executed on Sol 4. The checkout exercised science data collection and transfer, yet was a short duration such that if the anomalous behavior repeated, the system would not be overwhelmed with a large amount of data. The instrument checkout executed successfully on Sol 10 demonstrating that the anomalous behavior was transient.

*Root cause* – The investigation determined that the anomaly interrupted the flash read/write functions and that the effect was transient and could be cleared by power cycling the instrument. However, the mechanism that causes the flash read/write interruption is unknown. The PAE cannot be modified in flight, thus even if the root cause was determined to lie in PAE logic, no fix could be made. The InSight project, therefore, chose to focus on operational workarounds that would minimize instrument downtime and impacts to science and limit the effect of the anomaly on the rest of the lander system.

#### *Resuming Operations*

APSS returned to nominal operations on Sol 14. The anomaly investigation and recovery took nine days to complete, as shown in Table 1. APSS recovery activities also had to be planned around deployment activities, which contributed to the recovery time. If the anomaly had re-occurred, the overproduction of APSS data could have delayed the downlink of higher priority data and thus delayed the deployment timeline.

While APSS returned to operations, two changes were made to the lander's data handling processes in order to limit the impact of an overproduction of APSS data. First, the Special EVR buffer size was reduced. This lowered the number of Special EVR packets that would be downlinked following an anomaly. Secondly, APSS continuous science data was moved lower in the downlink data priority. This was done to reduce the possibility that a re-occurrence of the anomaly could delay the return of high-priority deployment data.

Sol Number	Activity
5	Anomaly occurs APSS powered off
8	Retrieve data from PAE
10	Repeat instrument checkout
14	Resume nominal instrument operations

**Table 1: Sol 5 Anomaly Investigation and Recovery Timeline**

#### *Subsequent Occurrences*

The Data Return anomaly has continued to occur throughout surface operations. The symptoms are the same each time, confirming that the same anomaly is occurring. Instances are random with no trends in time of sol, instrument on-time, flash memory utilization, or sleep/wake cycle as shown in Table 2 and Figure 6.


Sol Number	Instrument On-Time prior to Anomaly (sol)	Flash Memory Utilization	Sleep/Wake Cycle
5	0.2	14%	Sleep
30	7.2	51%	Sleep
92	16.3	55%	Wake
108	14.7	27%	Sleep
111	0.4	33%	Sleep
120	8.1	22%	Sleep
171	48.7	46%	Sleep
217	45.4	87%	Sleep
231	14.1	87%	Wake
260	28.0	54%	Wake
266	5.9	53%	Sleep
367	82.6	54%	Sleep
370	3.5	70%	Sleep

**Table 2: Anomaly Occurrences**

## **4. RECOVERY PROCESS**

To recover APSS following a data return anomaly, the instrument must be powered cycle. The instrument commands are simple enough; however, the complexity is introduced due to the operations planning cycle, staffing, and uplink opportunities. At first, the recovery process was tied to the operations shift schedule. Recovery commands were sent as part of the daily activity planning. Shifts were scheduled near daily, so this method allowed a turnaround time of approximately 48 hours, driven by shift time, uplink time, and onboard execution time. (With the exception of the first anomaly recovery which took nine days.)

As the end of the Deployment phase and daily operations shifts approached, the need for a more streamlined recovery process was realized. Once the transition to a weekly planning cycle occurred, the ability to respond to an anomaly via the nominal planning cycle would significantly hamper instrument operations. Instrument downtimes of more than a week could be possible depending on when an anomaly occurred. Thus, two alternative recovery approaches were proposed: 1) create an “on-the-shelf” command product, known as a load-and-go sequence, that could be uplinked as soon as possible after the ground received notification that the anomaly occurred and 2) create an onboard anomaly detection and recovery method that would remove the need for a ground-in-the-loop response. The InSight project elected to develop both methods. The first method would be the quickest to develop, test, and validate for inflight use thus serving as a stop-gap while the second, more complex, approach was developed.



**Figure 6: Time of Sol for each Anomaly Occurrence**

### *Load-and-Go Recovery Method*

The load-and-go (LGO) recovery method leveraged command products used during the daily planning cycles to recover the instrument. The commands are packaged as an LGO sequence, meaning that the commands execute as soon as they are uplinked to the lander. The advantage of this approach is that it was divorced from the daily or weekly planning cycle. The product could be developed and approved for use and be placed on-the-shelf until it was needed. However, ground operations team support would still be needed (to support product uplink<sup>3</sup>) in order to recover APSS.

*Constraints and Guidelines* – As LGO sequences execute as soon as they are received by the lander, the timing of sending such a sequence is very important and levied constraints on this recovery method. First, the recovery could not overlap with any other APSS commands as it could cause a command to not execute as intended and leave the instrument in an

unexpected state. This risk was addressed by unloading the APSS-dedicated sequence engines<sup>4</sup> prior to executing the instrument recovery. Second, the wake on which the LGO would be received by the lander had to be sufficiently long to execute the recovery. If the lander started a sleep shutdown before APSS commanding was completed, the instrument would be safed. This risk was mitigated by testing the LGO sequence thoroughly to confirm its execution duration and assessing the lander wake that would be used for the recovery and the uplink time of the sequence to ensure sufficient margin existed.

*LGO Logic* - The LGO sequence had three parts to it – the first part safed APSS in order to always start execution of the power cycle commands from a consistent instrument state. The delay between the anomaly occurrence and a ground response could result in the instrument safing due to other onboard activities. Thus, this step eliminated any uncertainties in the initial instrument state. The second part

<sup>3</sup> In a weekly plan, multiple uplink opportunities are available for contingency response use

<sup>4</sup> A sequencing engine is an FSW construct that dedicates computing space

to executing a sequence. InSight's FSW architecture is built to allocate specific engines to specific operations, as opposed to commands utilizing the first available engine.

unsafed the instrument, while the third part powered on the PAE and all the sensors. The use of this product was therefore limited to the nominal scenario of restoring all sensors to their powered state. At the time the LGO was developed, lander power resources were sufficient to allow all APSS sensors to be powered on, so this was an acceptable constraint.

*Verification and Validation* – While the commands had all been executed previously in flight, a verification and validation (V&V) effort still had to be carried out to approve the LGO sequence for inflight use. The V&V effort ensured that interactions between the commands were tested, the execution duration characterized, and all nominal and off-nominal execution scenarios evaluated based on different starting conditions, thus eliminating surprises due to unforeseen system-level responses. The V&V plan required the product to satisfy sequence construction rules, pass sequence validation checks, and be tested in high-fidelity system testbeds. These testbeds provided the most-flight like check as they model the interaction of commands with lander FSW and execute commands on the engineering model of the PAE.

A test plan was created which defined several cases, focusing on expected initial conditions, overlapping events and outcomes. Test cases included the instrument being safed or unsafed prior to the LGO's execution, instrument safing during the recovery, and LGO execution while the APSS sequence engine was already in use. All of these scenarios were modeled carefully in simulations. The telemetry and event records from commands were then analyzed to determine if any unexpected fault conditions occurred. Several iterations of the LGO product were generated based on the results of the testing. Some of the test cases were then repeated for completeness. The anomaly itself could not be replicated in any of the test venues, therefore the focus of the LGO sequence testing was on its functionality. Symptoms of the anomaly were not required to be replicated since operations in flight had proven the safe recovery of the instrument with a power cycle.

The V&V testing proved that the LGO command product successfully recovered the PAE in viable nominal and off-nominal scenarios.

*Results* – The LGO method improved recovery time because the command product was already available allowing the operations team to respond to an anomaly outside the nominal planning cycle. The LGO was successfully used in flight on four occasions with an average recovery time of 45 hours. The LGO was an important stop-gap capability between recovery via the daily tactical cycle and implementation of an automated method. All uses of this method occurred during the transition to the science monitoring phase when activity planning occurred only three times per week. Without the ability to uplink the recovery on non-planning days, instrument downtime would have doubled.

#### *Automated Detection and Recovery Method*

While the LGO method divorced instrument recovery from the nominal planning cycle, it still relied on operations team staffing. The drive was to find a recovery method that did not require actions from the operations team and would therefore minimize instrument downtime. However, to return the instrument to nominal operations without requiring a ground assessment meant occurrences of the anomaly had to be detectable by the lander. Thus, this auto-recovery method had to include both detection and recovery steps.

*Onboard Anomaly Detection* – In order to detect the anomaly onboard, the evidence of it must be visible to FSW. An overflow of the Special EVR buffer is detectable via onboard telemetry, however packets are only sent to this buffer after data is transferred from the PAE to the lander. Thus, the anomaly can only be detected when a science data processing session is executed during a full lander wake. The number of full lander wakes and science data processing sessions varies per sol, but the activity plan generally allows for one to three sessions per sol. The auto-recovery is designed to execute a check for an overflow of the Special EVR buffer directly after each session.

*Constraints and Guidelines* – The method had to consider timing constraints due to its dependence on available lander telemetry for detection, and the time needed to power cycle APSS for recovery. As discussed, the auto-recovery method can only be executed during a full lander wake. Also, sufficient time after a data transfer session must be guaranteed for a possible APSS power cycle – the lander must remain awake and no other APSS activities must be executed during the recovery. Operational guidelines were implemented to specify a minimum time that must be allocated to science data processing sessions and to prevent initiation of any other APSS activities until a power cycle would be complete.

As with the LGO method, the auto-recovery also needed to avoid sequence engine conflicts. By kicking off recovery automatically, there is some uncertainty as to what sequences may already be occupying the APSS and science data processing sequence engines. Thus, timing constraints were imposed on when the auto-recovery could run following a processing session and logic was built to ensure the recovery sequence would never collide with an already occupied engine.

Additionally, constraints were implemented on how the method executed instrument recovery. First, the auto-recovery cannot execute if APSS is already safed. This is to prevent the auto-recovery from responding to an unrelated anomaly. Second, limits were imposed on how often the auto-recovery can power cycle the instrument. If the response is run more than three times with less than 48 hours between each execution, the auto-recovery is aborted and APSS is safed. Frequent auto-recovery attempts could indicate an issue with the instrument and, thus, APSS operations are halted to allow the operations team to diagnose and respond.



Finally, the lander FSW does not have visibility into the power states of the APSS sensors (with the exception of TWINS). Thus, there is no direct way to determine the configuration of APSS at the time of the anomaly and return it to the same state which poses a risk that the instrument could be returned to a state that consumes more power than was planned. This risk was mitigated by using knowledge of the TWINS power state to determine which sensors should be powered on. If either one or both TWINS sensors are in use at the time of the anomaly, they are powered back on as is the IFG sensor. If neither TWINS boom is powered at the time of anomaly, then neither TWINS nor IFG are powered back on. The PS and 28V monitor are always powered back on by the auto-recovery due to their significance to science investigations and lander health.

*Auto-Recovery Logic* – The logic to automatically detect and recover from the APSS Data Return anomaly can be seen in Figure 7. The logic is contained in a reusable<sup>5</sup> sequence saved in lander memory and is called during a science data processing session. The ability of a sequence to check lander telemetry states and execute logic steps to determine the correct execution path was critical to the development of the auto-recovery.

The first step in the logic is to check that APSS is not already safed. If it is, no further action is taken and the response is aborted. Next, the lander telemetry is checked to determine if there is an overflow of packets from the Special EVR buffer. If this packet buffer overflow is detected, additional constraints are checked to ensure that auto-recovery is not occurring too frequently. To conduct these checks, two global variables were implemented: one global variable to track the number of power cycle attempts, and another to track the time of the most recent auto-recovery. If these checks pass, the response continues to instrument recovery. Otherwise, the response is aborted.

Auto-recovery first powers off the instrument. Then packets are cleared from the Special EVR buffer to remove bad packets resulting from the anomaly and also to make space in the buffer for the new packets that will be generated during instrument power on. Next, the PAE is powered on, as well as the 28V monitor and PS. Finally, the logic determines if at least one TWINS sensor was on prior to the anomaly. If so, the respective sensor(s) are powered on along with the IFG.

*Verification and Validation* – Implementing the automated response required the creation of new sequence and command products and modification of several existing command products. The test plan needed to individually test each product for command correctness, format, flight rule compliance, and also comprehensively for system-level interactions and responses. End-to-end system-level testing was especially significant given the automated nature of the onboard implementation.

At the unit level, each product was tested in a similar way as described previously for the LGO method. Scenarios were identified for each product's use in flight, and mapped with expected responses. These were modeled in the appropriate test venue and checked for correctness.

One of the key aspects of the design of the auto-recovery testing was simulating the anomaly in the test venues. This was applicable to the auto-recovery products because they relied on onboard detection of certain symptoms to take different actions. However, because the anomaly could not be recreated in the test venues, steps were taken to simulate the symptoms of the anomaly – in this case, the overflow of the Special EVR. By configuring the test venue's buffer space to be artificially small, the buffer overflow was simulated, which allowed the testing of the various logical paths in the command products.

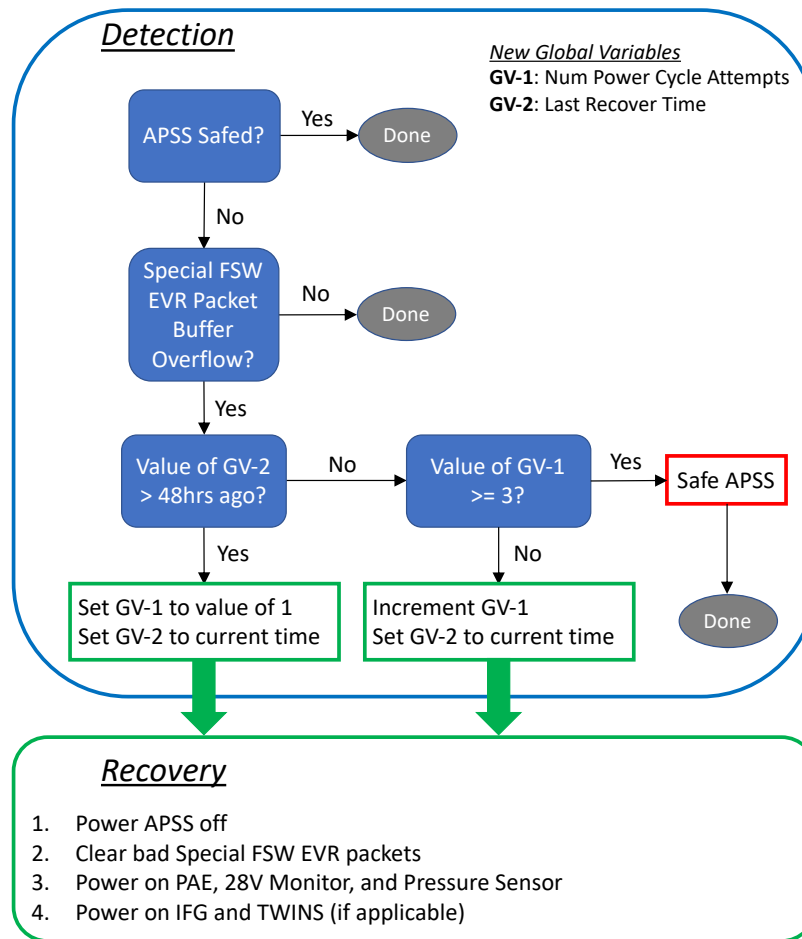
A comprehensive system-level plan was then designed. The primary goal was to execute the relevant command products in as flight-like a test environment as possible. The system plan modeled several lander wake and shutdown cycles and various combinations of test scenarios were included in each wake cycle, for example, different instrument power states and parallel lander activities (such as science data processing and communication windows). The test sequences stressed the timing and limits of the arguments used by the command products to identify race conditions and any unforeseen issues. All the system tests were executed in the high-fidelity system testbed which allowed testing the commands with real flight software interactions, and simulated accurate timing of events and computing loads.

Several iterations of testing allowed the operations team to update timing in some of the products in order to allow margin for all activities to complete in nominal or off-nominal situations. Testing confirmed that the products worked together as expected, telemetry and global variables updated appropriately, and that the auto-recovery was only executed within the limits defined by the method logic.

*Results* – The auto-recovery method is a marked improvement in responding to the APSS Data Return Anomaly. This approach has removed the need to stop instrument operations and wait for a ground assessment and response. The detection and recovery from the anomaly can all be accomplished via FSW and by leveraging sequence architecture that allows sequences to be stored onboard the lander and to evaluate telemetry states and determine the execution path. The auto-recovery has been used numerous times in flight and the average instrument downtime is nine hours. This is an order of five improvement in instrument downtime compared to the LGO method.

<sup>5</sup> Sequences usually delete themselves at execution time. By omitting this self-delete, the sequence remains in lander memory and can be executed

multiple times.



**Figure 7: Auto-Recovery Logic**

#### *Updates to the Automated Detection and Recovery Method*

The InSight project pursued an update to the SEIS/APSS FSW (for an unrelated issue), which opened the door to other minor change requests. Improving system visibility into the state of APSS would allow further efficiencies to be incorporated into the auto-recovery method. The project agreed to an update such that APSS RTHK packets are read by the lander and select status values, including the Flash Inhibit flag and sensor power states, are placed into a lander global variable. This global variable is then utilized by auto-recovery.

**New Detection Method** – By detecting an anomaly via the Flash Inhibit flag, the auto-recovery performs faster, more frequent checks for the anomaly. The anomaly can now be detected at any lander wake, as APSS RTHK packets are transferred during both diagnostic and full lander wakes. There is no longer a constraint to wait for a science data processing session to complete during a full lander wake before being able to detect an occurrence of the anomaly. Leveraging all lander wakes further minimizes instrument downtime and science data loss.

**Minimization of System-Level Impacts** – In this updated method, if the Flash Inhibit flag is detected then no science data processing will occur, thus preventing corrupted, unusable data from being sent to the lander and ground and therefore mitigating buffer overflow, inflated downlink data volumes, and overwhelming ground processing pipelines. By utilizing the known state of the instrument sensors, the auto-recovery determines more accurately and quickly which sensors to repower following an anomaly. This removes the risk of returning APSS to a greater power-consuming state than originally planned.

**Reduction of Auto-Recovery Execution Time** – The auto-recovery sequence was updated to minimize execution time. First, when APSS is unpowered, the power switch to the PAE is commanded open rather than powering off all the sensors first, as was originally done in the auto-recovery. There is no harm in removing power to the PAE without powering off the sensors first, all instruments are designed to handle such an occurrence.

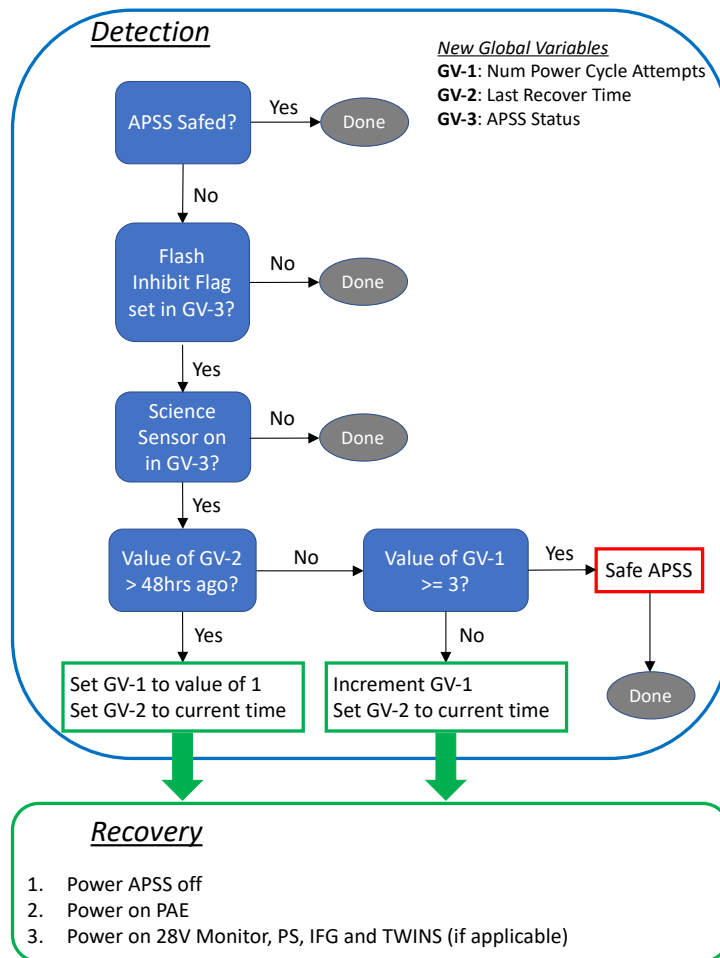
Second, when returning to operations, auto-recovery originally leveraged existing command products to power on

each sensor. However, these command products required waiting at least one minute between successive sensor power on commands (due to PAE restrictions) which added significant time to the recovery process. The updated auto-recovery determines which sensors to repower and builds a single command to repower all the selected sensors, thus requiring only a single minute wait following this command. While this implementation saves substantial execution time for the power-cycle sequence, it utilizes a “pass-thru” command for which the command parameters are built within the sequence itself. This is not a typical means to command the instrument and required project approval as well as substantial V&V testing to verify that the sequence (and associated pass-thru command) worked as expected.

By enabling the updated auto-recovery to execute more rapidly, the detection and response now completes prior to the start of instrument activities during any lander wake. Thus, the nominal activity plan will no longer be interrupted due to this anomaly. Also, operational planning complexity is reduced because additional time no longer needs to be allocated after a science data processing session for a possible execution of auto-recovery. This means other lander activities, such as starting a sleep cycle, can occur without risk of safing APSS because it is still executing an activity.

*Updated Auto-Recovery Logic* – The updated auto-recovery logic is shown in Figure 8. In this approach, an anomaly occurrence is checked for at the start of every lander wakeup and prior to performing any instrument data transfers. First, the APSS safed state is checked. If the instrument is already safed no further actions are taken. If APSS is not safed and the Flash Inhibit flag is set, then the auto-recovery determines the appropriate response: do nothing if no science sensors are powered, safe the instrument if the response is executing too frequently, or proceed to recovering the instrument which will power cycle the instrument back to the same state.

*Results* – The additional visibility into APSS status and reduced runtime of the updated auto-recovery method enables occurrences of the anomaly to be detected more quickly. The method modifications also reduced the system-level impacts of the anomaly. The modified SEIS/APSS FSW and updated auto-recovery method were implemented onboard the lander in November 2019. Over two subsequent anomaly occurrences, the updated auto-recovery method resulted in an instrument downtime reduction of 3.5 hours on average.



**Figure 8: Updated Auto-Recovery Logic**

## 5. LESSONS LEARNED

The telemetry and sequence capabilities of the InSight lander were essential to implementing an automated detection and response method. However, limitations in each of these capabilities, also introduced complexity to the method. Both supply lessons learned that could be applied to future mission operations.

### *Onboard Telemetry Visibility*

The ability to see telemetry values onboard and use them in sequence and command execution provides a powerful tool for real-time knowledge of the system state. For auto-recovery, being able to see the telemetry that indicated an overflow of the Special EVR buffer had occurred was essential to detecting and responding to the anomaly onboard. However, this onboard visibility only extended to lander telemetry and did not provide any insight into the contents of instrument science or engineering packets. The lack of visibility into APSS RTHK packets limited the knowledge of the instrument state and introduced some risk into the recovery method (this limitation was addressed by the SEIS/APSS FSW update). Future mission engineers should consider operational scenarios where onboard telemetry visibility would be useful and also trade what types of telemetry (e.g. spacecraft vs. payload) are valuable.

### *Multiple Sequence Execution Paths*

The sequence architecture used by InSight allowed the auto-recovery to leverage evaluation (if/else) statements and choose an execution path based on the current lander and instrument states. This room for uncertainty enabled a response that could be run from any initial instrument state. If a fixed, known instrument state would have been required to recover from the anomaly, it may have confined the operations team to a ground-in-the-loop response. While multiple sequence execution paths did require a more complex V&V campaign, in the case of the auto-recovery, it allowed for a robust series of checks and balances that were paramount to automatically responding to the anomaly. Future mission engineers should consider how flexibility in sequencing can enable operations and reduce brittleness to changes in instrument or system state.

## 6. SUMMARY

The APSS Data Return anomaly is a persistent anomaly with the potential to greatly impact instrument operations and science data return. The automated detection and response method leverages the sequence and telemetry capabilities of the lander to implement a unique approach that significantly reduces instrument downtime (see Figure 9) and eliminates ground-in-the-loop interaction. The efficacy of this approach

has been proven several times in flight. Most significantly during an occurrence just prior to the August 2019 conjunction period<sup>6</sup>, where, without the automated detection and response method, APSS data collection may have been halted approximately four weeks until the operations team could respond.


## ACKNOWLEDGEMENTS

This work was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

## REFERENCES

- [1] Smrekar, S.E., Lognonné, P., Spohn, T. et al. Space Sci Rev (2019) 215: 3. <https://doi.org/10.1007/s11214-018-0563-9>
- [2] Lognonné, P., Banerdt, W.B., Giardini, D. et al. Space Sci Rev (2019) 215: 12. <https://doi.org/10.1007/s11214-018-0574-6>
- [3] Banfield, D., Rodriguez-Manfredi, J.A., Russell, C.T. et al. Space Sci Rev (2019) 215: 4. <https://doi.org/10.1007/s11214-018-0570-x>
- [4] Spohn, T., Grott, M., Smrekar, S.E. et al. Space Sci Rev (2018) 214: 96. <https://doi.org/10.1007/s11214-018-0531-4>
- [5] Folkner, W.M., Dehant, V., Le Maistre, S. et al. Space Sci Rev (2018) 214: 100. <https://doi.org/10.1007/s11214-018-0530-5>
- [6] Trebi-Ollennu, A., Kim, W., Ali, K. et al. Space Sci Rev (2018) 214: 93. <https://doi.org/10.1007/s11214-018-0520-7>
- [7] Hoffman T., "InSight: Mission to Mars," 2018 IEEE Aerospace Conference, Big Sky, MT, 2018, pp. 1-11. <https://doi.org/10.1109/AERO.2018.8396723>
- [8] InSight APSS Critical Design Review, April 10, 2014
- [9] InSight Payload Auxiliary Electronics FPGA Design Document, May 8, 2015
- [10] InSight PAE Critical Design Review, May 8, 2014
- [11] SEIS FSW User's Manual, April 14, 2016
- [12] InSight APSS Operations Manual, March 7, 2019

<sup>6</sup> Conjunction occurs when the Sun-Earth-Mars angle is significantly low and Earth-Mars communications are compromised



**Figure 9: Improvement in Instrument Downtime over Evolution of Anomaly Response**

#### BIOGRAPHY



*Emily Manor-Chapman is the APSS Instrument Engineer for the InSight Mission. She has previously worked on the Cassini Mission, Europa Clipper Pre-Project, and the Juno Mission. Currently, she is the Deployment Phase Lead for the NASA-ISRO Synthetic Aperture Radar Mission. She has a BS in Aerospace Engineering and a MEng in Space Systems Engineering from the University of Michigan.*



*Elizabeth Barrett is the Science and Instrument Operations Team Lead for the InSight mission. She has previously worked on the Jason-3 mission at JPL and has eight years of operational experience at Johnson Space Center as a flight controller for the International Space Station. She has a BA (astronomy) and MEng (engineering physics) from Cornell University and a PhD in astronomy from University of Hawaii.*



**Farah Alibay** obtained a PhD from the Massachusetts Institute of Technology in 2014 and has been a systems engineer at the Jet Propulsion Laboratory since then. She has worked on a variety of Mars missions. On InSight, she was the payload V&V engineer during development, and she worked as a Tactical Uplink shift

Lead (TUL) during deployment and commissioning operations.



**Kyle Cloutier** is a Sequencing Integration Engineer for the InSight Mission. She has previously worked on the Mars Exploration Rovers (MER) Mission and the Cassini Mission, and is currently a Payload Systems Engineer with the Surface Water and Ocean Topography (SWOT) and Sentinel-6 Missions.

She has a BS in Aerospace Engineering from the University of Maryland.



**Jonathan Grinblat** is Lead Payload Systems Engineer of the InSight project at the Jet Propulsion Laboratory, California Institute of Technology. He has worked on the successful Mars Exploration Rover and Mars Science Laboratory projects, as well technology development programs. Specialties include Payload Systems

Engineering, Avionics Systems Engineering, and Digital Electronics Design, and Integration and Test.



**Jesse Mendoza Jr.** is currently a Mission System Verification & Validation system engineer for the Mars 2020 rover mission at NASA's Jet Propulsion Laboratory. Since graduating from the University of California Riverside in 2017, Jesse has done verification & validation work at JPL for the Orbiting Carbon Observatory 3 (OCO-3) and Mars

InSight missions, as well as mission operations for InSight. Jesse's interests include space systems engineering, space science, emerging technologies, public outreach, music, and overall space exploration.



**Nimisha Mittal** is the Instrument command product V&V lead on InSight. She has previously worked on the operations teams for the Mars Exploration rovers and Cassini missions, and payload systems on the Mars Science Laboratory mission. She holds a BS in Aerospace Engineering from the

University of Texas at Austin and an MS from Stanford University. Currently she is working as a flight operations development and V&V engineer on the NISAR mission, a collaboration between NASA and ISRO.